

ArbiLoMod, a Simulation Technique Designed for Arbitrary Local Modifications

Andreas Buhr ^{†‡§} Christian Engwer [‡] Mario Ohlberger [‡]
 Stephan Rave ^{‡¶}

November 30, 2016

Engineers manually optimizing a structure using Finite Element based simulation software often employ an iterative approach where in each iteration they change the structure slightly and resimulate. Standard Finite Element based simulation software is usually not well suited for this workflow, as it restarts in each iteration, even for tiny changes. In settings with complex local microstructure, where a fine mesh is required to capture the geometric detail, localized model reduction can improve this workflow. To this end, we introduce ArbiLoMod, a method which allows fast recomputation after arbitrary local modifications. It employs a domain decomposition and a localized form of the Reduced Basis Method for model order reduction. It assumes that the reduced basis on many of the unchanged domains can be reused after a localized change. The reduced model is adapted when necessary, steered by a localized error indicator. The global error introduced by the model order reduction is controlled by a robust and efficient localized a posteriori error estimator, certifying the quality of the result. We demonstrate ArbiLoMod for a coercive, parameterized example with changing structure.

Key words. model order reduction, reduced basis method, domain decomposition, a posteriori error estimation

AMS subject classifications. 65N55, 65N30

1 Introduction

Finite Element based simulation is a standard tool in many CAD/CAE assisted workflows in engineering. Depending on the complexity of the design simulated, on the underlying partial differential equation, and on the desired fidelity of the approximation of the solution, performing a simulation may take hours, days or even weeks.

ArbiLoMod aims at the acceleration of a very specific class of problems, namely the repetitive simulation of parameterized problems with fine microstructure without scale separation: Problems which exhibit a microstructure on a scale much smaller than the domain require a very fine mesh to resolve the geometry and thus take very long to compute. They often have much more degrees of freedom than necessary for the description

[†]Corresponding author. ✉ andreas@andreasbuhr.de

[‡]Institute for Computational and Applied Mathematics, University of Münster, Einsteinstraße 62, 48149 Münster, Germany

[§]supported by CST Computer Simulation Technology AG

[¶]supported by the German Federal Ministry of Education and Research (BMBF) under contract number 05M13PMA




Figure 1.1: DDR memory channel on a printed circuit board subject to local modification of conductive tracks.

of their physical behavior, thus model order reduction can succeed. For problems with scale separation, there are established methods to reduce the model such as MsFEM [27], HMM [17], VMM [28], or LOD [38]. However, when there is no scale separation, homogenization is not possible. ArbiLoMod is based on a localized RB approach which does not require scale separation.

In addition, engineers often want to change the problem definition and resimulate several times in an iterative process. Their goal is not to come from a problem definition to an approximation of the solution, but to provide approximations to the solutions of a sequence of problems, where each problem was created by modifying the previous problem. We assume that changes might be arbitrary, i.e. non parametric, but are of local nature, so that the majority of the problem setting (e.g. large parts of the geometry) is unchanged between two simulation runs. ArbiLoMod exploits the similarities between subsequent problems by reusing local approximation spaces in regions of the computation domain unaffected by the change.

Furthermore, for many design applications in each iteration, the structure under consideration has to be simulated for a multitude of model parameters to analyze its behavior. ArbiLoMod includes online offline decomposition techniques from RB methods [23, 42, 21], using the regularity structure of the solution manifold [6] for fast many-query simulation of the model.

A particular example that we have in mind is the design of printed circuit boards (PCBs). The design of PCBs has all of the above mentioned properties: PCBs are nowadays very complex, there is no scale separation, improvements are often obtained by local changes of the electronic components and conductive tracks, and when solved in frequency domain, it is a parameterized problem with the frequency as a parameter. A possible change is depicted in Figure 1.1. The same applies to integrated circuit (IC) packages, which are in structure similar to PCBs. Other areas of application could be e.g. resonance analysis in cars or trains or electromagnetic filter design.

ArbiLoMod's goal is the reduction of overall simulation times. When measuring simulation times, we assume that the runtime on a single workstation is not the right quantity to look at. Instead, the runtime on any hardware easily available to the user is what matters, which particularly includes massive computing power in cluster and cloud environments, but usually not millions of cores as in supercomputers. As the user still has to pay per compute node in cloud environments, hundreds to thousands of compute nodes is the foreseen environment. Secondary goals during the development of ArbiLoMod were that the method should be easily implementable on top of existing finite element schemes and that the detection of changed regions vs. unchanged regions between two simulation runs is completely automatic.

After an overview over existing methods in the literature, the definition of the problem setting and a short overview of ArbiLoMod, the structure of the paper follows the structure of ArbiLoMod: In Section 3 the space decomposition used in this paper is presented. Training and Greedy algorithms for local basis generation are subject of Section 4. The a posteriori error estimator employed is discussed Section 5. Localized enrichment of the bases is described in Section 6. The procedure followed on each geometry change is given in Section 7. Potentials for parallelization are sketched in Section 8. Section 9 contains

numerical results.

Existing Approaches

The combination of ideas of the fields of Reduced Basis Methods, of multiscale methods and domain decomposition methods gained a lot of attention in recent time. In 2002, Maday and Rønquist published the “Reduced Basis Element” (RBE) method [33, 34], combining the reduced basis approach with a domain decomposition, coupling local basis vectors by polynomial Lagrange multipliers on domain boundaries. Built on top of the RBE is the “Reduced Basis Hybrid Method” (RBHM) [31] and the “Discontinuous Galerkin Reduced Basis Element” method (DGRBE) [4]. Similar in motivation is the “Static Condensation Reduced Basis Element” (SCRBE) method [41], which also aims at systems composed of components, where the geometry of the components can be mapped to reference geometries. While the connection between the components is simply achieved by polynomial Lagrange multipliers in the RBE, significant research has been conducted on choosing the right coupling spaces in the context of the SCRBE. Choosing the right space at the interfaces is called “Port Reduction” and was included in the name, leading to the “Port Reduced Static Condensation Reduced Basis Element” (PR-SCRBE) method [19, 20] which employs a so-called “pairwise training”. A variation of this idea is used in our method. Recently, also an algorithm to obtain optimal interface spaces for the PR-SCRBE was proposed [44] and a framework for a posteriori error estimation was introduced [43]. While the PR-SCRBE performs excellent in using the potentials of cloud environments, its goals are different from ours. While being able to handle various changes to the system simulated, it does not aim at arbitrary modifications. And it does not try to hide the localization from the user, but rather exposes it to let the user decide how the domain decomposition should be done. Another combination of domain decomposition ideas with reduced basis methods coupling different physical formulations on the domain boundaries was presented in [35, 36].

A completely different approach to local model order reduction is to see it as an extension to multiscale methods. There are two ways to combine RB and multiscale methods. First is to use RB to accelerate the solution of localized problems which occur in multiscale methods (“RB within multiscale”). This has been done by multiple authors, see e.g. [1, 2, 24]. The second way is to use a subspace projection for the global problem, but use ideas from multiscale methods to construct the basis functions. This second approach is used by ArbiLoMod and is shared with several methods in the literature: The “Generalized Multiscale Finite Element Method” GMsFEM [18] uses the idea of “Multiscale Finite Elements” MsFEM [27] and constructs reduced spaces which are spanned by ansatz functions on local patches, using only local information. It allows for non-fixed number of ansatz functions on each local patch. GMsFEM uses local eigenproblems and a partition of unity for basis generation. Adaptive enrichment for the GMsFEM is presented by Chung et al. in [14, 12] and online-adaptive enrichment in [13]. While an application of the GMsFEM to the problem of arbitrary local modifications would be very interesting, GMsFEM was not designed to be communication avoiding. A parallel implementation of the enrichment described in [13] would require the communication of high dimensional basis representations, which is avoided in ArbiLoMod.

There are also developments of the “Generalized Finite Element Method” GFEM [46, 47] which could be extended to handle arbitrary local modifications. Especially the recent development of the Multiscale-GFEM [5] could be promising in this regard.

Similar to ArbiLoMod in spirit is the “Localized Reduced Basis Multiscale Method” (LRBMS) [3, 39]. It uses a non overlapping domain decomposition and discontinuous ansatz spaces, which are coupled using a DG ansatz at the interface. While LRBMS could be extended to handle arbitrary local modifications and it can be implemented in a communication avoiding scheme, LRBMS cannot easily be implemented on top of an existing

conforming discretization scheme. ArbiLoMod, in contrast, inherits all conformity properties from the underlying discretization. It can be built on top of standard, conforming finite element schemes and is then conforming itself. Some of the basic ideas of ArbiLoMod were already published by the authors [7, 10]. First results for electrodynamics were published in [9].

The main advantage of ArbiLoMod is its speed in cloud environments. At every design decision during the development of ArbiLoMod, care was taken to keep the required communication in a parallel implementation at a minimum.

2 Preliminaries

Problem Setting

In this contribution we particularly look at problems that are modeled by partial differential equations with complex local structure, typically on a very fine scale compared to the overall problem setting. In addition, we assume that the problem might depend on a number of parameters $\mu \in \mathcal{P} \subset \mathbb{R}^p$. As a model problem to explore and design our new simulation technique ArbiLoMod, we consider elliptic equations with complex micro-structure. To simplify the presentation, we restrict ourselves to the two dimensional case in this publication. Thus, let $\Omega \subset \mathbb{R}^d$, $d = 2$ denote the polygonal computational domain, and V , $H_0^1(\Omega) \subset V \subset H^1(\Omega)$ denote the solution space. We then look at variational problems of the form

$$\text{find } u_\mu \in V \text{ such that } a_\mu(u_\mu, v) = \langle f_\mu, v \rangle \quad \forall v \in V. \quad (2.1)$$

Here, $a_\mu : V \times V \rightarrow \mathbb{R}$, $\mu \in \mathcal{P}$ denotes a parameterized bilinear form with inherent micro-structure and $f_\mu \in V'$ a force term. V is equipped with the standard H^1 inner product and the thereby induced norm. a_μ is assumed to be coercive and continuous and by α_μ, γ_μ we denote the lower (upper) bounds for the coercivity (norm) constants of a_μ , i.e. $\alpha_\mu \|\varphi\|_V^2 \leq a_\mu(\varphi, \varphi)$ for all $\varphi \in V$ and $\|a_\mu\| \leq \gamma_\mu$. We further assume that the bilinear form $a_\mu(u, v)$ has a decomposition affine in the parameters and can be written as a sum of parameter independent bilinear forms $a^b(u, v)$ with parameter dependent coefficient functions $\theta_b(\mu)$:

$$a_\mu(u, v) = \sum_b \theta_b(\mu) a^b(u, v) \quad (2.2)$$

As an example a_μ could be given as

$$a_\mu(u, v) = \int_\Omega \sigma_\mu(x) \nabla u(x) \nabla v(x) dx \quad (2.3)$$

where $\sigma_\mu : \Omega \rightarrow \mathbb{R}$ denotes a parameterized heat conduction coefficient that varies in space on a much finer scale than the length scale of Ω .

Structure of ArbiLoMod

The main ingredients of ArbiLoMod are:

1. a localizing space decomposition,
2. local training and greedy algorithms,
3. a localized a posteriori error estimator,
4. a localized adaptive enrichment procedure.

ArbiLoMod builds upon a space decomposition of the original ansatz space V consisting of a set of subspaces $V_i \subset V$ whose direct sum is the original ansatz space, i.e.

$$V = \bigoplus_i V_i. \quad (2.4)$$




Figure 2.1: Overview of ArbiLoMod. Generation of initial reduced interface spaces by training and of reduced volume spaces by greedy basis generation is subject of Section 4. Convergence is assessed with the localized a posteriori error estimator presented in Section 5. Enrichment is discussed in Section 6. On each geometry change, the procedure starts over, where new initial interface spaces are only generated in the region affected by the change.

The subspaces V_i are meant to have local properties, i.e. all elements of one subspace V_i have support only in a small subset of the domain. In the first step, for each local subspace V_i , a reduced local subspace $\tilde{V}_i \subseteq V_i$ is constructed using training and greedy algorithms. Thereafter, the global problem is solved in the space formed by the direct sum of all reduced local subspaces:

$$\begin{aligned}
 &\text{with} \quad \tilde{V} := \bigoplus_i \tilde{V}_i \\
 &\text{find} \quad \tilde{u}_\mu \in \tilde{V} \quad \text{such that} \quad a_\mu(\tilde{u}_\mu, v) = \langle f_\mu, v \rangle \quad \forall v \in \tilde{V}.
 \end{aligned} \tag{2.5}$$

To assess the quality of the thus obtained solution, a localized a posteriori error estimator is employed. If necessary, the solution is improved by enriching the reduced local subspaces, using a residual based, localized enrichment procedure. Finally, on each localized change, affected bases are discarded and the procedure starts over. An overview is given in Figure 2.1.

3 Space Decomposition

In this section we introduce definitions for the different kinds of spaces needed in the method. Since ArbiLoMod is designed to work on top of existing discretization schemes, we formulate the method on a discrete level. The following definitions are for two dimensional problems, but can be easily extended for three dimensional problems.

3.1 Basic Subspaces

V_h denotes a discrete ansatz space, spanned by ansatz functions $\{\psi_i\}_{i=1}^N =: \mathcal{B}$. We assume that the ansatz functions have a localized support, which is true for many classes of ansatz functions like Lagrange- or Nédélec-type functions. We first introduce a direct decomposition of the ansatz space V_h into “basic subspaces”. These are used to construct the space decomposition later. To obtain the subspaces we classify the ansatz functions by their support and define each subspace as the span of all ansatz functions of one class. To this end, we introduce a non overlapping domain decomposition of the original domain

Ω into open subdomains Ω_i :

$$\overline{\Omega} = \bigcup_{i=1}^{N_D} \overline{\Omega}_i \quad \Omega_i \cap \Omega_j = \emptyset \text{ for } i \neq j \quad (3.1)$$

where N_D is the number of subdomains. For each ψ in \mathcal{B} , we call I_ψ the set of indices of subdomains that have non-empty intersection with the support of ψ , i.e.

$$I_\psi := \left\{ i \in \{1, \dots, N_D\} \mid \text{supp}(\psi) \cap \Omega_i \neq \emptyset \right\}. \quad (3.2)$$

We collect all occurring domain sets in Υ :

$$\Upsilon := \left\{ I_\psi, \psi \in \mathcal{B} \right\}. \quad (3.3)$$

We define sets for each codimension:

$$\Upsilon_0 := \left\{ \xi \in \Upsilon \mid |\xi| = 1 \right\}, \quad \Upsilon_1 := \left\{ \xi \in \Upsilon \mid |\xi| = 2 \right\}, \quad \Upsilon_2 := \left\{ \xi \in \Upsilon \mid |\xi| > 2 \right\}. \quad (3.4)$$

The elements of Υ_2, Υ_1 , and Υ_0 can be associated with interior vertices (codimension 2), faces (codimension 1) and cells (codimension 0) of the domain decomposition. The classification is very similar to the classification of mesh nodes in domain decomposition methods, see for example [32, Def. 3.1] or [48, Def. 4.2].

Definition 3.1 (Basic Subspaces) For each element $\xi \in \Upsilon$ we define a basic subspace U_ξ of V as:

$$U_\xi := \text{span} \left\{ \psi \in \mathcal{B} \mid I_\psi = \xi \right\}.$$

Remark 3.2 (Basic Decomposition) The definition of U_ξ induces a direct decomposition of V_h :

$$V_h = \bigoplus_{\xi \in \Upsilon} U_\xi.$$

3.2 Space Decomposition

As mentioned in the introduction, ArbiLoMod is based on a space decomposition. It can work on the Basic Decomposition introduced in Definition 3.1. However, faster convergence and smaller basis sizes are achieved using the modified space decomposition introduced in this section. Here we assume that the discrete ansatz space V_h is spanned by finite element ansatz functions defined on a mesh which resolves the subdomains Ω_i .

For each of the spaces U_ξ defined in Definition 3.1, we calculate extensions. The extensions are computed on the “extension space” $E(U_\xi)$ which is defined as

$$E(U_\xi) := \bigoplus \left\{ U_\zeta \mid \zeta \subseteq \xi \right\}. \quad (3.5)$$

Examples for extension spaces are given in Figure 3.1. For each space U_ξ , a linear extension operator Extend is defined:

$$\text{Extend} : U_\xi \rightarrow E(U_\xi). \quad (3.6)$$

For all ξ in Υ_0 , Extend is just the identity. For all ξ in Υ_1 , we extend by solving the homogeneous version of the equation with Dirichlet zero boundary values for one (arbitrary) chosen $\bar{\mu} \in \mathcal{P}$. For example, in the situation depicted in Fig. 3.2b,

$$\begin{aligned} \text{Extend} : U_{\{1,2\}} &\rightarrow U_{\{1\}} \oplus U_{\{1,2\}} \oplus U_{\{2\}} \\ \varphi &\mapsto \varphi + \psi \\ &\text{where } \psi \in U_{\{1\}} \oplus U_{\{2\}} \text{ solves} \\ &a_{\bar{\mu}}(\varphi + \psi, \phi) = 0 \quad \forall \phi \in U_{\{1\}} \oplus U_{\{2\}} \end{aligned}$$




Figure 3.1: Visualization of basic spaces U_ξ and their extension spaces for Q^1 ansatz functions (one dof per mesh node).




Figure 3.2: Visualization of some example elements of the local subspaces V_ξ for inhomogeneous coefficients. The structure in the solution results from variations in the heat conduction coefficient.

For all ξ in Υ_2 , *Extend* is defined by first extending linearly to zero on all edges in the extension domain, i.e. in all spaces in $E(U_\xi)$ which belong to Υ_1 . Then, in a second step, the homogeneous version of the equation with Dirichlet boundary values is solved on the spaces in $E(U_\xi)$ which belong to Υ_0 . The procedure is visualized in Figure 3.3. The functions constructed by this two step procedure are the same as the MsFEM basis functions used by Hou and Wu [27]. Note that the base functions for $E(U_\xi)$, $\xi \in \Upsilon_2$ form a partition of unity in the interior of the coarse partition of the domain. They can be completed to form a partition of unity on the whole coarse partition of the domain if suitable base functions for the vertices at the boundary of the domain are added. This will be used in Section 5 below for the robust and efficient localization of an a posteriori error estimator. Examples of extended functions for all codimensions are given in Figure 3.2. In the case of the Laplace equation these basis functions coincide with the hat functions on the coarse partition (see Figure 3.3). As discussed in Section 5 below, the choice of hat function can be an alternative choice that allows for a better a priori bound of the constants in the localized a posteriori error estimator, as their gradient is controlled by $1/H$ – where H denotes the mesh size of the macro partition – independent of the contrast of the data.

For the communication avoiding properties of the ArbiLoMod, it is important to note that extensions can be calculated independently on each domain, i.e. $\text{Extend}(\varphi)|_{\Omega_i}$ can be calculated having only information about φ and Ω_i , without knowledge about other domains. Using this operator, we define the local subspaces V_ξ :

Definition 3.3 (Extended Subspaces) *For each element $\xi \in \Upsilon$ we define an extended subspace V_ξ of V_h as:*

$$V_\xi := \left\{ \text{Extend}(\varphi) \mid \varphi \in U_\xi \right\}.$$

According to the definition in (3.4) we call V_ξ a cell, face, or vertex space, if $\xi \in \Upsilon_0$, $\xi \in \Upsilon_1$,




Figure 3.3: Extend operator is executed in two steps for spaces $V_{\xi}, \xi \in \Upsilon_2$. Exemplified for the generation of $V_{\{1,2,3,4\}}$ from $U_{\{1,2,3,4\}}$. Mesh and spaces as depicted in Fig 3.1. Script: `generate_vertex_extension_plots.nb`.

or $\xi \in \Upsilon_2$, respectively (cf. Fig. 3.2).

Remark 3.4 (Extended Decomposition) *The definition of V_{ξ} induces a direct decomposition of V_h :*

$$V_h = \bigoplus_{\xi \in \Upsilon} V_{\xi}.$$

Space decompositions of the same spirit are used in the context of Component Mode Synthesis (CMS), see [26, 25].

3.3 Projections

Definition 3.5 (Local Projection Operators) *The projection operators $P_{U_{\xi}} : V_h \rightarrow U_{\xi}$ and $P_{V_{\xi}} : V_h \rightarrow V_{\xi}$ are defined by the relation*

$$\varphi = \sum_{\xi \in \Upsilon} P_{U_{\xi}}(\varphi) = \sum_{\xi \in \Upsilon} P_{V_{\xi}}(\varphi) \quad \forall \varphi \in V_h.$$

As both the subspaces U_{ξ} and the subspaces V_{ξ} form a direct decomposition of the space V_h , the projection operators are uniquely defined by this relation.

The implementation of the projection operators $P_{U_{\xi}}$ is very easy: It is just extracting the coefficients of the basis functions forming U_{ξ} out of the global coefficient vector. The implementation of the projection operators $P_{V_{\xi}}$ is more complicated and involves the solution of local problems, see Algorithm 1.

Algorithm 1: Projections in V_{ξ}

```

1 Function SpaceDecomposition( $\varphi$ ):
   Input : function  $\varphi \in V_h$ 
   Output: decomposition of  $\varphi$ 
2   /* iterate over all codimensions in decreasing order */
3   for codim  $\in \{d, \dots, 0\}$  do
4       for  $\xi \in \Upsilon_{\text{codim}}$  do
5            $\varphi_{\xi} \leftarrow \text{Extend}(P_{U_{\xi}}(\varphi))$ 
6            $\varphi \leftarrow \varphi - \varphi_{\xi}$ 
7   return  $\{\varphi_{\xi}\}$ 
    
```

4 Local Basis Generation

For each local subspace V_ξ , an initial reduced local subspace $\tilde{V}_\xi \subseteq V_\xi$ is generated, using only local information from an environment around the support of the elements in U_ξ . The strategy used to construct these reduced local subspaces depends on the type of the space: whether ξ belongs to Υ_0 , Υ_1 or Υ_2 . The three strategies are given in the following. The local basis generation algorithms can be run in parallel, completely independent of each other. See Section 8 for further discussion of the potential parallelization. As the algorithms only use local information, their results do not change when the problem definition is changed outside of the area they took into account. So there is no need to rerun the algorithms in this case. Our numerical results indicate that the spaces obtained by local trainings and greedys have good approximation properties (see Section 9). The quality of the obtained solution will be guaranteed by the a posteriori error estimator presented in Section 5 below.

4.1 Basis Construction for Reduced Vertex Spaces

The spaces V_ξ for $\xi \in \Upsilon_2$ are spanned by only one function (see Figure 3.2c for an example) and are thus one dimensional. The reduced spaces are therefore chosen to coincide with the original space, i.e. $\tilde{V}_\xi := V_\xi, \forall \xi \in \Upsilon_2$.

4.2 Local Training for Basis Construction of Reduced Face Spaces

To generate an initial reduced local subspace for V_ξ , $\xi \in \Upsilon_1$ we use a local training procedure. Its main four steps are to

1. solve the equation on a small domain around the space in question with zero boundary values for all parameters in the training set Ξ ,
2. solve the homogeneous equation repeatedly on a small domain around the space in question with random boundary values for all parameters in Ξ ,
3. apply the space decomposition to all obtained local solutions to obtain the part belonging to the space in question and
4. use a greedy procedure to create a space approximating this set.

The complete algorithm is given in Algorithm 3 and explained below.

The training is inspired by the “Empirical Port Reduction” introduced in Eftang et al. [19] but differs in some key points. The main differences are: (1) Within [19], the trace of solutions at the interface to be trained is used. This leads to the requirement that interfaces between domains do not intersect. In ArbiLoMod, a space decomposition is used instead. This allows ports to intersect, which in turn allows the decomposition of space into domains. (2) The “Empirical Port Reduction” trains with a pair of domains. We use an environment of the interface in question, which contains six domains in the 2D case. In 3D, it contains 18 domains. (3) PR-SCRBE aims at providing a library of domains which can be connected at their interfaces. The reduced interface spaces are used in different domain configurations and have to be valid in all of them. Within the context of ArbiLoMod, no database of domains is created and the interface space is constructed only for the configuration at hand, which simplifies the procedure. (4) The random boundary values used in [19] are generalized Legendre polynomials with random coefficients. In ArbiLoMod, the finite element basis functions with random coefficients are used, which simplifies the construction greatly, especially when there is complex structure within the interface.

The Training Space

To train a basis for \tilde{V}_ξ , $\xi \in \Upsilon_1$, we start from the subspace U_ξ . For each subspace U_ξ , we define a corresponding training space $T(U_\xi)$ on an environment associated with the




Figure 4.1: Visualization of basic spaces $U_{\{2,5\}}$, its training space, and the coupling space of its training space for Q^1 ansatz functions (one dof per mesh node).

face ξ . The following definition is geometric and tailored to a domain decomposition in rectangular domains. More complex domain decompositions would need a more complex definition here. We define the neighborhood

$$\mathcal{N}_{\xi} := \left\{ i \in \{1, \dots, N_D\} \mid \overline{\Omega}_i \cap \left(\bigcap_{k \in \xi} \overline{\Omega}_k \right) \neq \emptyset \right\} \quad (4.1)$$

and with that the training spaces

$$T(U_{\xi}) := \bigoplus \left\{ U_{\zeta} \mid \zeta \subseteq \mathcal{N}_{\xi} \right\}. \quad (4.2)$$

The training space is coupled to the rest of the system via its coupling space

$$C(T(U_{\xi})) := \bigoplus \left\{ U_{\zeta} \mid \zeta \cap \mathcal{N}_{\xi} \neq \emptyset, \zeta \not\subseteq \mathcal{N}_{\xi} \right\}. \quad (4.3)$$

A sketch of the degrees of freedom associated with the respective spaces is given in Figure 4.1. These definitions are also suitable in the 3D case. We have fixed the size of the neighborhood to one domain from the interface in question in each direction. This facilitates the setup of local problems and the handling of local changes: After a local change, the affected domains are determined. Afterwards, all trainings have to be redone for those spaces which contain an affected domain in their training domain. While a larger or smaller training domain might be desirable in some cases (see [22]), it is not necessary: As missing global information is added in the enrichment step, ArbiLoMod always converges to the desired accuracy, even if the training domain is not of optimal size. So the advantages of having the size of the training domain fixed to one domain outweighs its drawbacks.

The reduced basis must be rich enough to handle two types of right hand sides up to a given accuracy $\varepsilon_{\text{train}}$: (a) source terms and boundary conditions, and (b) arbitrary values on the coupling interface, both in the whole parameter space \mathcal{P} . We define an extended parameter space $\mathcal{P} \times C(T(U_{\xi}))$. For this parameter space we construct a training space $\Xi \times G \subset \mathcal{P} \times C(T(U_{\xi}))$, where G denotes an appropriate sampling of $C(T(U_{\xi}))$. We use the finite element basis $\mathcal{B}_{C(T(U_{\xi}))}$ on the coupling space and generates M random coefficient vectors r_i of size $N_B = \dim(C(T(U_{\xi})))$. With this an individual coupling space function $\varphi \in C(T(U_{\xi}))$ is constructed as

$$\varphi_i = \sum_{j=1}^{N_B} r_{ij} \phi_j; \quad \phi_j \in \mathcal{B}_{C(T(U_{\xi}))}. \quad (4.4)$$

For our numerical experiments in Section 9, we use uniformly distributed random coefficients over the interval $[-1, 1]$ and Lagrange basis functions. For each $\mu \in \Xi$ and each pair

$(\mu, g_c) \in \Xi \times G$ we construct snapshots u_f and u_c as solutions for right hand sides $f_\mu(\cdot)$, $a_\mu(g_c, \cdot)$ respectively, i.e.

$$\begin{aligned} a_\mu(u_f, \phi) &= \langle f_\mu, \phi \rangle \quad \forall \phi \in T(U_\xi), \\ a_\mu(u_c, \phi) &= -a_\mu(g_c, \phi) \quad \forall \phi \in T(U_\xi). \end{aligned} \quad (4.5)$$

Based on the set of snapshots which we call Z , a reduced basis \mathcal{B} is constructed using a greedy algorithm (cf. Alg. 2). In the numerical experiments, the V -norm and V -inner product are used. The complete generation of the reduced face spaces \tilde{V}_ξ , $\xi \in \Upsilon_1$ with basis $\tilde{\mathcal{B}}_{\tilde{V}_\xi}$ is summarized in Alg. 3.

Algorithm 2: SnapshotGreedy

```

1 Function SnapshotGreedy( $Z, \varepsilon_{\text{train}}$ ):
    Input : set of elements to approximate  $Z$ ,
           training tolerance  $\varepsilon_{\text{train}}$ 
    Output: basis of approximation space  $\mathcal{B}$ 
2  $\mathcal{B} \leftarrow \emptyset$ 
3 while  $\max_{z \in Z} \|z\| > \varepsilon_{\text{train}}$  do
4      $\hat{z} \leftarrow \arg \max_{z \in Z} \|z\|$ 
5      $\hat{z} \leftarrow \frac{\hat{z}}{\|\hat{z}\|}$ 
6      $Z \leftarrow \{z - (z, \hat{z})\hat{z} \mid z \in Z\}$ 
7      $\mathcal{B} \leftarrow \mathcal{B} \cup \{\hat{z}\}$ 
8 return  $\mathcal{B}$ 
    
```

Algorithm 3: Training to construct reduced face spaces \tilde{V}_ξ , $\xi \in \Upsilon_1$

```

1 Function Training( $\xi, M, \varepsilon_{\text{train}}$ ):
    Input : space identifier  $\xi$ ,
           number of random samples  $M$ ,
           training tolerance  $\varepsilon_{\text{train}}$ 
    Output: reduced local subspace  $\tilde{V}_\xi$ 
2  $G \leftarrow \text{RandomSampling}(\xi, M)$ 
3  $Z \leftarrow \emptyset$ 
4 foreach  $\mu \in \Xi$  do
5     find  $u_f \in T(U_\xi)$  such that:
6          $a_\mu(u_f, \phi) = f_\mu(\phi) \quad \forall \phi \in T(U_\xi)$ 
7      $Z \leftarrow Z \cup P_{V_\xi}(u_f)$ 
8     foreach  $g_c \in G$  do
9         find  $u_c \in T(U_\xi)$  such that:
10             $a_\mu(u_c + g_c, \phi) = 0 \quad \forall \phi \in T(U_\xi)$ 
11             $Z \leftarrow Z \cup P_{V_\xi}(u_c)$ 
12  $\tilde{\mathcal{B}}_{\tilde{V}_\xi} \leftarrow \text{SnapshotGreedy}(Z, \varepsilon_{\text{train}})$ 
13 return  $\text{span}(\tilde{\mathcal{B}}_{\tilde{V}_\xi})$ 
    
```

4.3 Basis Construction for Reduced Cell Spaces Using Local Greedy

For each cell space V_ξ , $\xi \in \Upsilon_0$ we create a reduced space \tilde{V}_ξ . These spaces should be able to approximate the solution in the associated part of the space decomposition for any

variation of functions from reduced vertex or face spaces that are coupled with it. We define the reduced coupling space $\tilde{C}(V_\xi)$ and its basis $\tilde{\mathcal{B}}_{\tilde{C}(V_\xi)}$.

$$\Upsilon_\xi^C := \left\{ \zeta \in \Upsilon \mid \zeta \cap \xi \neq \emptyset, \zeta \not\subseteq \xi \right\} \quad (4.6)$$

$$\tilde{C}(V_\xi) := \bigoplus_{\zeta \in \Upsilon_\xi^C} \tilde{V}_\zeta \quad \tilde{\mathcal{B}}_{\tilde{C}(V_\xi)} := \bigcup_{\zeta \in \Upsilon_\xi^C} \tilde{\mathcal{B}}_{\tilde{V}_\zeta} \quad (4.7)$$

We introduce an extended training set: $\Xi \times \{1, \dots, N_{\tilde{B}} + 1\}$, $N_{\tilde{B}} := \dim(\tilde{C}(V_\xi))$. Given a pair $(\mu, j) \in \Xi \times \{1, \dots, N_{\tilde{B}} + 1\}$, we define the associated right hand side as

$$g_{\mu,j}(\phi) := \begin{cases} -a_\mu(\psi_j, \phi) & \text{if } j \leq N_{\tilde{B}} \\ \langle f_\mu, \phi \rangle & \text{if } j = N_{\tilde{B}} + 1, \end{cases} \quad (4.8)$$

where ψ_j denotes the j -th basis function of $\tilde{\mathcal{B}}_{\tilde{C}(V_\xi)}$. We then construct the reduced cell space \tilde{V}_ξ as the classical reduced basis space (cf. Alg. 4, LocalGreedy) with respect to the following parameterized local problem: Given a pair $(\mu, j) \in \Xi \times \{1, \dots, N_{\tilde{B}} + 1\}$, find $u_{\mu,j} \in V_\xi$ such that

$$a_\mu(u_{\mu,j}, \phi) = g_{\mu,j}(\phi) \quad \forall \phi \in V_\xi. \quad (4.9)$$

The corresponding reduced solutions are hence defined as: Find $\tilde{u}_{\mu,j} \in \tilde{V}_\xi$ such that:

$$a_\mu(\tilde{u}_{\mu,j}, \phi) = g_{\mu,j}(\phi) \quad \forall \phi \in \tilde{V}_\xi \quad (4.10)$$

Both problems have unique solutions due to the coercivity and continuity of a_μ . For the LocalGreedy (Algorithm 4) we use the standard Reduced Basis residual error estimator, i.e.

$$\|u_{\mu,j} - \tilde{u}_{\mu,j}\|_{V_\xi} \leq \Delta_{cell}(\tilde{u}_{\mu,j}) := \frac{1}{\alpha_{LB}(\mu)} \|R_{\mu,j}(\tilde{u}_{\mu,j})\|_{V'_\xi}, \quad (4.11)$$

with the local residual

$$\begin{aligned} R_{\mu,j} : V_\xi &\rightarrow V'_\xi \\ \varphi &\mapsto g_{\mu,j}(\cdot) - a_\mu(\varphi, \cdot) \end{aligned} \quad (4.12)$$

and a lower bound for the coercivity constant α_{LB} . The idea of using a local greedy to generate a local space for all possible boundary values can also be found in [30, 4].

5 A-Posteriori Error Estimator

The model reduction error in the ArbiLoMod has to be controlled. To this end, an a posteriori error estimator is used which should have the following properties:

1. It is robust and efficient.
2. It is online-offline decomposable.
3. It is parallelizable with little amount of communication.
4. After a localized geometry change, the offline computed data in unaffected regions can be reused.
5. It can be used to steer adaptive enrichment of the reduced local subspaces.

All these requirements are fulfilled by the estimator presented in the following. We develop localized bounds for the standard RB error estimator,

$$\Delta(\tilde{u}_\mu) := \frac{1}{\alpha_\mu} \|R_\mu(\tilde{u}_\mu)\|_{V'_h} \quad (5.1)$$

where $R_\mu(\tilde{u}_\mu) \in V'_h$ is the global residual given as $\langle R_\mu(\tilde{u}_\mu), \varphi \rangle = \langle f_\mu, \varphi \rangle - a_\mu(\tilde{u}_\mu, \varphi)$. This error estimator is known to be robust and efficient ([23, Proposition 4.4]):

$$\|u_\mu - \tilde{u}_\mu\|_V \leq \Delta(\tilde{u}_\mu) \leq \frac{\gamma_\mu}{\alpha_\mu} \|u_\mu - \tilde{u}_\mu\|_V. \quad (5.2)$$

Algorithm 4: LocalGreedy to construct local cell spaces \tilde{V}_ξ , $\xi \in \Upsilon_0$

```

1 Function LocalGreedy( $\xi, \varepsilon_{\text{greedy}}$ ):
    Input : space identifier  $\xi$ ,
            greedy tolerance  $\varepsilon_{\text{greedy}}$ 
    Output: reduced local subspace  $\tilde{V}_\xi$ 
2  $\tilde{B}_{\tilde{V}_\xi} \leftarrow \emptyset$ 
3 while  $\max_{\substack{\mu \in \Xi \\ j \in \{1, \dots, N_{\tilde{B}}+1\}}} \Delta_{\text{cell}}(\tilde{u}_{\mu,j}) > \varepsilon_{\text{greedy}}$  do
4      $\hat{\mu}, \hat{j} \leftarrow \arg \max_{\substack{\mu \in \Xi \\ j \in \{1, \dots, N_{\tilde{B}}+1\}}} \Delta_{\text{cell}}(\tilde{u}_{\mu,j})$ 
5     find  $u_{\hat{\mu}, \hat{j}} \in V_\xi$  such that:
6      $\lfloor a_{\hat{\mu}}(u_{\hat{\mu}, \hat{j}}, \phi) = g_{\hat{\mu}, \hat{j}}(\phi) \quad \forall \phi \in V_\xi$ 
7      $u_{\hat{\mu}, \hat{j}} \leftarrow u_{\hat{\mu}, \hat{j}} - \sum_{\phi \in \tilde{B}_{\tilde{V}_\xi}} (\phi, u_{\hat{\mu}, \hat{j}})_V \phi$ 
8      $\tilde{B}_{\tilde{V}_\xi} \leftarrow \tilde{B}_{\tilde{V}_\xi} \cup \{\|u_{\hat{\mu}, \hat{j}}\|_V^{-1} u_{\hat{\mu}, \hat{j}}\}$ 
9 return span( $\tilde{B}_{\tilde{V}_\xi}$ )
    
```

5.1 Abstract Estimates

We start by showing two abstract localized estimates for the dual norm of a linear functional.

Proposition 5.1 *Let $\{O_\xi\}_{\xi \in \Upsilon_E}$ be a collection of linear subspaces of V_h for some finite index set Υ_E and let $\tilde{V} \subset V_h$ denote an arbitrary subspace. Moreover let $P_{O_\xi} : V_h \rightarrow O_\xi \subseteq V_h$, $\xi \in \Upsilon_E$ be continuous linear mappings which satisfy $\sum_{\xi \in \Upsilon_E} P_{O_\xi} = \text{id}_{V_h}$. With the stability constant of this partition modulo \tilde{V} defined as*

$$c_{\text{pu}, \tilde{V}} := \sup_{\varphi \in V_h \setminus \{0\}} \frac{(\sum_{\xi \in \Upsilon_E} \inf_{\tilde{\varphi} \in \tilde{V} \cap O_\xi} \|P_{O_\xi}(\varphi) - \tilde{\varphi}\|_V^2)^{\frac{1}{2}}}{\|\varphi\|_V},$$

we have for any linear functional $f \in V'_h$ with $\langle f, \tilde{\varphi} \rangle = 0 \quad \forall \tilde{\varphi} \in \tilde{V}$ the estimate

$$\|f\|_{V'_h} \leq c_{\text{pu}, \tilde{V}} \cdot \left(\sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi}^2 \right)^{\frac{1}{2}},$$

where $\|f\|_{O'_\xi}$ denotes the norm of the restriction of f to O_ξ .

Proof. Using the Cauchy-Schwarz inequality and $\langle f, \tilde{\varphi} \rangle = 0 \quad \forall \tilde{\varphi} \in \tilde{V}$, we have

$$\begin{aligned}
 \|f\|_{V'_h} &= \sup_{\varphi \in V_h \setminus \{0\}} \frac{\sum_{\xi \in \Upsilon_E} \langle f, P_{O_\xi}(\varphi) \rangle}{\|\varphi\|_V} = \sup_{\varphi \in V_h \setminus \{0\}} \frac{\sum_{\xi \in \Upsilon_E} \inf_{\tilde{\varphi} \in \tilde{V} \cap O_\xi} \langle f, P_{O_\xi}(\varphi) - \tilde{\varphi} \rangle}{\|\varphi\|_V} \\
 &\leq \sup_{\varphi \in V_h \setminus \{0\}} \frac{\sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi} \inf_{\tilde{\varphi} \in \tilde{V} \cap O_\xi} \|P_{O_\xi}(\varphi) - \tilde{\varphi}\|_V}{\|\varphi\|_V} \\
 &\leq \sup_{\varphi \in V_h \setminus \{0\}} \frac{(\sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi}^2)^{\frac{1}{2}} (\sum_{\xi \in \Upsilon_E} \inf_{\tilde{\varphi} \in \tilde{V} \cap O_\xi} \|P_{O_\xi}(\varphi) - \tilde{\varphi}\|_V^2)^{\frac{1}{2}}}{\|\varphi\|_V} \\
 &= c_{\text{pu}, \tilde{V}} \cdot \left(\sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi}^2 \right)^{\frac{1}{2}}.
 \end{aligned}$$

□

A stability constant very similar to $c_{\text{pu},\tilde{V}}$ appears in the analysis of overlapping domain decomposition methods (e.g. [48, Assumption 2.2], [45]) and in localization of error estimators on stars (e.g. [15]).

Proposition 5.2 *With the assumptions in Proposition 5.1, let $\dot{\bigcup}_{j=1}^J \Upsilon_{E,j} = \Upsilon_E$ be a partition of Υ_E such that*

$$\forall 1 \leq j \leq J \ \forall \xi_1 \neq \xi_2 \in \Upsilon_{E,j} : O_{\xi_1} \perp O_{\xi_2}.$$

Then we have

$$\left(\sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi}^2 \right)^{\frac{1}{2}} \leq \sqrt{J} \|f\|_{V'_h},$$

Proof. Let $O_{\xi_1} \perp O_{\xi_2}$ be some subspaces of V_h , and let f be a continuous linear functional on $O_{\xi_1} \oplus O_{\xi_2}$. If $v_{f,1} \in O_{\xi_1}$ and $v_{f,2} \in O_{\xi_2}$ are the Riesz representatives of the restrictions of f to O_{ξ_1} and O_{ξ_2} , then due to the orthogonality of O_{ξ_1} and O_{ξ_2} , $v_{f,1} + v_{f,2}$ is the Riesz representative of f on $O_{\xi_1} \oplus O_{\xi_2}$. Thus,

$$\begin{aligned} \|f\|_{(O_{\xi_1} \oplus O_{\xi_2})'}^2 &= \|v_{f,1} + v_{f,2}\|_{V_h}^2 \\ &= \|v_{f,1}\|_{V_h}^2 + \|v_{f,2}\|_{V_h}^2 = \|f\|_{O'_{\xi_1}}^2 + \|f\|_{O'_{\xi_2}}^2, \end{aligned}$$

where we have used the orthogonality of the spaces again. The same is true for a larger orthogonal sum of spaces. We therefore obtain:

$$\begin{aligned} \sum_{\xi \in \Upsilon_E} \|f\|_{O'_\xi}^2 &= \sum_{j=1}^J \sum_{\xi \in \Upsilon_{E,j}} \|f\|_{O'_\xi}^2 \\ &= \sum_{j=1}^J \|f\|_{(\oplus_{\xi \in \Upsilon_{E,j}} O_\xi)'}^2 \\ &\leq J \|f\|_{V'_h}^2. \end{aligned}$$

□

When grouping the spaces O_ξ so that in each group, all spaces are orthogonal to each other, J is the number of groups needed. Applying both estimates to the residual, we obtain an efficient, localized error estimator:

Corollary 5.3 *The error estimator $\Delta_{loc}(\tilde{u}_\mu)$ defined as*

$$\Delta_{loc}(\tilde{u}_\mu) := \frac{1}{\alpha_\mu} c_{\text{pu},\tilde{V}} \left(\sum_{\xi \in \Upsilon_E} \|R_\mu(\tilde{u}_\mu)\|_{O'_\xi}^2 \right)^{\frac{1}{2}} \quad (5.3)$$

is robust and efficient:

$$\|u_\mu - \tilde{u}_\mu\|_V \leq \Delta_{loc}(\tilde{u}_\mu) \leq \frac{\gamma_\mu \sqrt{J} c_{\text{pu},\tilde{V}}}{\alpha_\mu} \|u_\mu - \tilde{u}_\mu\|_V$$

Proof. Applying Propositions 5.1 and 5.2 to the error estimator $\Delta(\tilde{u}_\mu) = \frac{1}{\alpha_\mu} \|R_\mu(\tilde{u}_\mu)\|_{V'_h}$ yields, together with (5.2), the proposition. □

Online-offline decomposition of this error estimator can be done by applying the usual strategy for online-offline decomposition used with the standard RB error estimator (see e.g. [23, Sec. 4.2.5] or the numerically more stable approach [8]) to every dual norm in $\Delta_{loc}(\tilde{u}_\mu)$.

5.2 Choosing Spaces

The error estimator defined in Corollary 5.3 works for any spaces O_ξ and mappings P_{O_ξ} fulfilling the assumptions in Proposition 5.1. However, in order to obtain good constants $c_{\text{pu},\tilde{V}}$ and \sqrt{J} , both have to be chosen carefully. In addition, two more properties are needed for good performance of the implementation. First, the subspaces should be spanned by FE ansatz functions, allowing the residual to be easily evaluated on these spaces. Second, the inner product matrix on the subspaces should be sparse, as the inner product matrix has to be solved in the computation of the dual norms. We use an overlapping decomposition based on the non-overlapping domain decomposition introduced in (3.1).

Definition 5.4 (Overlapping space decomposition) *Let the index set Υ_E for the overlapping space decomposition be given by the vertices of the domain decomposition, i.e.*

$$\Upsilon_E = \Upsilon_2. \quad (5.4)$$

We then define the overlapping spaces O_ξ supported on the overlapping domains Ω_ξ by:

$$O_\xi := \bigoplus \left\{ U_\zeta \mid \zeta \subseteq \xi \right\}, \quad \Omega_\xi := \bigcup_{i \in \xi} \overset{\circ}{\Omega}_i \quad \xi \in \Upsilon_E. \quad (5.5)$$

Note that we have $O_\xi = \{\psi \in \mathcal{B} \mid \text{supp}(\psi) \subseteq \overset{\circ}{\Omega}_\xi\} \subseteq H_0^1(\Omega_\xi)$. Contrary to V_ξ or U_ξ , these spaces do not form a direct sum decomposition of V_h . We next state a first estimate on the partition of unity constant of Corollary (5.3) for this choice of partition, which does not take into account that the residual vanishes on the reduced space. The resulting estimate thus depends on H^{-1} , where H is the minimum diameter of the subdomains of the macro partition. Typically the size of the macro partition is moderate such that H^{-1} is small. However, in the following Proposition 5.7 we will show that the constant $c_{\text{pu},\tilde{V}}$ can be actually bounded independent of H , when we choose a partition of unity that is contained in the reduced space \tilde{V} .

Proposition 5.5 *Let $c_{\text{ovlp}} := \max_{x \in \Omega} \#\{\xi \in \Upsilon_E \mid x \in \Omega_\xi\}$ be the maximum number of estimator domains Ω_ξ overlapping in any point x of Ω and let $H_\xi := \text{diam}(\Omega_\xi)$, $\xi \in \Upsilon_E$ and $H := \min_{\xi \in \Upsilon_E} H_\xi$. Furthermore, assume that there exist partition of unity functions $p_\xi \in H^{1,\infty}(\Omega)$, $\xi \in \Upsilon_E$ and a linear interpolation operator $\mathcal{I} : V \rightarrow V_h$ such that*

- (i) $\sum_{\xi \in \Upsilon_E} p_\xi(x) = 1$ for all $x \in \Omega$,
- (ii) $\max_{\xi \in \Upsilon_E} \|p_\xi\|_\infty \leq 1$ and $\|\nabla p_\xi\|_\infty \leq c'_{\text{pu}} H_\xi^{-1}$ for all $\xi \in \Upsilon_E$,
- (iii) $\mathcal{I}(\varphi) = \varphi$ for all $\varphi \in V_h$,
- (iv) $\mathcal{I}(p_\xi V_h) \subseteq O_\xi$ for all $\xi \in \Upsilon_E$,
- (v) $\|\mathcal{I}(p_\xi v_h) - p_\xi v_h\|_V \leq c_I \|v_h\|_{\Omega_\xi,1}$ for all $\xi \in \Upsilon_E, v_h \in V_h$.

Then we have:

$$c_{\text{pu},\tilde{V}} \leq \sqrt{4 + 2c_I^2 + 4(c'_{\text{pu}} H^{-1})^2} \cdot \sqrt{c_{\text{ovlp}}}.$$

Proof. We compute the bound for c_{pu} using the partition of unity and the interpolation operator. To this end, let

$$P_{O_\xi}(\varphi) := \mathcal{I}(p_\xi \varphi), \quad \xi \in \Upsilon_E.$$

Due to (iv), these are linear mappings $V_h \rightarrow O_\xi$, and using (i) and (iii) we obtain $\sum_{\xi \in \Upsilon_E} P_{O_\xi}(\varphi) = \mathcal{I}(\sum_{\xi \in \Upsilon_E} p_\xi \varphi) = \mathcal{I}(\varphi) = \varphi$ for all $\varphi \in V_h$. Thus, Corollary 5.3 applies with this specific choice of partition operators P_{O_ξ} . Now, using (ii) and (v) we have for

any $\varphi \in V_h$

$$\begin{aligned}
 \sum_{\xi \in \Upsilon_E} \|P_{O_\xi}(\varphi)\|_V^2 &\leq 2 \sum_{\xi \in \Upsilon_E} \|\mathcal{I}(p_\xi \varphi) - p_\xi \varphi\|_V^2 + \|p_\xi \varphi\|_V^2 \\
 &\leq 2 \sum_{\xi \in \Upsilon_E} c_I^2 \|\varphi\|_{\Omega_{\xi,1}}^2 + \int_{\Omega_\xi} 2|\nabla p_\xi \varphi|^2(x) + 2|p_\xi \nabla \varphi|^2(x) + |p_\xi \varphi|^2(x) dx \\
 &\leq 2 \sum_{\xi \in \Upsilon_E} c_I^2 \|\varphi\|_{\Omega_{\xi,1}}^2 + (1 + 2(c'_{pu} H^{-1})^2) |\varphi|_{\Omega_{\xi,0}}^2 + 2|\varphi|_{\Omega_{\xi,1}}^2 \\
 &\leq (4 + 2c_I^2 + 4(c'_{pu} H^{-1})^2) c_{\text{ovlp}} \|\varphi\|_V^2
 \end{aligned}$$

This gives us the estimate. \square

Remark 5.6 When the domain decomposition Ω_i is sufficiently regular (e.g. see the numerical examples below), partition of unity functions satisfying (i) and (ii) can easily be found. If $V_h \cup \{p_\xi \mid \xi \in \Upsilon_E\}$ consists of p -th order finite element basis functions for some fine triangulation of Ω , Lagrange interpolation can be chosen as interpolation operator \mathcal{I} . In fact, using standard interpolation error estimates and inverse inequalities one sees that for each element T of the fine triangulation with diameter h one has:

$$\begin{aligned}
 \|\mathcal{I}(p_\xi v_h) - p_\xi v_h\|_{T,1} &\leq ch^p |p_\xi v_h|_{T,p+1} \\
 &\leq c' h^p \sum_{k=1}^p |v_h|_{T,k} \cdot |p_\xi|_{T,p+1-k,\infty} \\
 &\leq c'' h^p \sum_{k=1}^p h^{-(k-1)} |v_h|_{T,1} \cdot h^{-(p+1-k)} |p_\xi|_{T,0,\infty} \\
 &\leq c'' p |v_h|_{T,1},
 \end{aligned}$$

where c'' is a constant bounded by the shape regularity of the fine triangulation.

For the rectangular domain decomposition used in the numerical example below, the constant J is $J = 2^d = 4$: it is possible to divide the overlapping domains into four classes, so that within each class, no domain overlaps with any other (cf. [13, Sec. 5]).

Furthermore, the coercivity constant α_μ and the stability constant γ_μ , or estimates, are required. For the numerical example presented in Section 9, those can be calculated analytically. In general this is not possible and the details of estimating them numerically are subject for further investigations. The numerical computation of a lower bound for the coercivity constant was subject of extensive research in the RB community (see e.g. [29, 11]), but these methods require the calculation of the coercivity constant at some parameter values and thus require the solution of a global, fine scale problem. To the authors' knowledge, there are no publications on localization of these methods so far.

The upper bound on the constant $c_{pu,\tilde{V}}$ in Proposition 5.5 depends on the domain size H approximately like H^{-1} . As the domain size is considered a constant in the context of ArbiLoMod, the error estimator is already considered efficient with this bound. In the next proposition, we however show that the constant can indeed be bounded independent of H , if we exploit that the residual vanishes on the reduced space \tilde{V} .

Proposition 5.7 Let $p_\xi, \xi \in \Upsilon_E$ be a partition of unity and \mathcal{I} an interpolation operator satisfying the prerequisites of Proposition 5.5. Furthermore, assume $V = H_0^1(\Omega)$ and that $p_\xi \in \tilde{V} \cap O_\xi$ for $\xi \in \Upsilon_E^{\text{int}} := \{\xi \in \Upsilon_E \mid \bar{\Omega}_\xi \cap \partial\Omega = \emptyset\}$, e.g. p_ξ is chosen as a basis function of \tilde{V}_ξ (see Subsections 3.2 and 4.1 above). Then the following estimate holds:

$$c_{pu,\tilde{V}} \leq \sqrt{4 + 2c_I^2 + 4(c'_{pu} c_{pc})^2} \cdot \sqrt{c_{\text{ovlp}}},$$

with a Poincaré-inequality constant c_{pc} (see proof below) that does not depend on the fine or coarse mesh sizes (h, H) .

Proof. For arbitrary $\varphi \in V_h$ let $\bar{\varphi}_\xi := \frac{1}{|\Omega_\xi|} \int_{\Omega_\xi} \varphi$. We then have with $\Upsilon_E^{\text{ext}} := \Upsilon_E \setminus \Upsilon_E^{\text{int}}$

$$\begin{aligned} c_{\text{pu}, \tilde{V}} &= \sup_{\varphi \in V_h \setminus \{0\}} \frac{(\sum_{\xi \in \Upsilon_E} \inf_{\tilde{\varphi} \in \tilde{V} \cap O_\xi} \|P_{O_\xi}(\varphi) - \tilde{\varphi}\|_V^2)^{\frac{1}{2}}}{\|\varphi\|_V} \\ &\leq \sup_{\varphi \in V_h \setminus \{0\}} \frac{(\sum_{\xi \in \Upsilon_E^{\text{int}}} \|P_{O_\xi}(\varphi) - \bar{\varphi}_\xi p_\xi\|_V^2 + \sum_{\xi \in \Upsilon_E^{\text{ext}}} \|P_{O_\xi}(\varphi)\|_V^2)^{\frac{1}{2}}}{\|\varphi\|_V}, \end{aligned}$$

where we have used that by construction $\bar{\varphi}_\xi p_\xi \in \tilde{V} \cap O_\xi$ for all $\xi \in \Upsilon_E^{\text{int}}$. For any $\varphi \in V_h$ and $\xi \in \Upsilon_E^{\text{int}}$ we then have $\|P_{O_\xi}(\varphi) - \bar{\varphi}_\xi p_\xi\|_V^2 \leq 2c_I^2 \|\varphi\|_{\Omega_{\xi,1}}^2 + 2\|(\varphi - \bar{\varphi}_\xi)p_\xi\|_V^2$, where

$$\begin{aligned} \|(\varphi - \bar{\varphi}_\xi)p_\xi\|_V^2 &\leq \int_{\Omega_\xi} 2|\nabla(\varphi - \bar{\varphi}_\xi)p_\xi|^2(x) + 2|(\varphi - \bar{\varphi}_\xi)\nabla p_\xi|^2(x) dx \\ &\quad + \|(\varphi - \bar{\varphi}_\xi)p_\xi\|_{L^2(\Omega_\xi)}^2. \end{aligned}$$

With a rescaled Poincaré-type inequality

$$\|\varphi - \bar{\varphi}_\xi\|_{L^2(\Omega_\xi)} \leq c_{\text{pc}} H_\xi \|\nabla \varphi\|_{L^2(\Omega_\xi)},$$

and $\|\varphi - \bar{\varphi}_\xi\|_{L^2(\Omega_\xi)} \leq \|\varphi\|_{L^2(\Omega_\xi)}$, we get

$$\begin{aligned} \int_{\Omega_\xi} 2|\nabla(\varphi - \bar{\varphi}_\xi)p_\xi|^2(x) + 2|(\varphi - \bar{\varphi}_\xi)\nabla p_\xi|^2(x) dx + \|(\varphi - \bar{\varphi}_\xi)p_\xi\|_{L^2(\Omega_\xi)}^2 \\ \leq (2 + 2(c'_{\text{pu}} c_{\text{pc}})^2) \|\nabla \varphi\|_{L^2(\Omega_\xi)}^2 + \|\varphi\|_{L^2(\Omega_\xi)}^2 \end{aligned}$$

In analogy we obtain for the boundary terms, i.e. $\xi \in \Upsilon_E^{\text{ext}}$, the estimates $\|P_{O_\xi}(\varphi)\|_V^2 \leq 2c_I^2 \|\varphi\|_{\Omega_{\xi,1}}^2 + 2\|\varphi p_\xi\|_V^2$, and

$$\begin{aligned} \|\varphi p_\xi\|_V^2 &\leq \int_{\Omega_\xi} 2|\nabla \varphi p_\xi|^2(x) + 2|\varphi \nabla p_\xi|^2(x) dx + \|\varphi p_\xi\|_{L^2(\Omega_\xi)}^2 \\ &\leq (2 + 2(c'_{\text{pu}} c_{\text{pc}})^2) \|\nabla \varphi\|_{L^2(\Omega_\xi)}^2 + \|\varphi\|_{L^2(\Omega_\xi)}^2. \end{aligned}$$

using a rescaled Poincaré-type inequality which holds for $\xi \in \Upsilon_E^{\text{ext}}$ as $\varphi \in V_h$ has zero boundary values, i.e.

$$\|\varphi\|_{L^2(\Omega_\xi)} \leq c_{\text{pc}} H_\xi \|\nabla \varphi\|_{L^2(\Omega_\xi)}.$$

Summing up all contributions we then have

$$\begin{aligned} \sum_{\xi \in \Upsilon_E^{\text{int}}} \|P_{O_\xi}(\varphi) - \bar{\varphi}_\xi p_\xi\|_V^2 + \sum_{\xi \in \Upsilon_E^{\text{ext}}} \|P_{O_\xi}(\varphi)\|_V^2 \\ \leq \sum_{\xi \in \Upsilon} 2c_I^2 \|\varphi\|_{\Omega_{\xi,1}}^2 + 2 \left[(2 + 2(c'_{\text{pu}} c_{\text{pc}})^2) \|\nabla \varphi\|_{L^2(\Omega_\xi)}^2 + \|\varphi\|_{L^2(\Omega_\xi)}^2 \right] \\ \leq (4 + 2c_I^2 + 4(c'_{\text{pu}} c_{\text{pc}})^2) c_{\text{ovlp}} \|\varphi\|_V^2. \end{aligned}$$

This gives us the estimate. \square

Proposition 5.7 gives a bound on $c_{\text{pu}, \tilde{V}}$ that depends on the contrast of the underlying diffusion coefficient if $p_\xi \in \tilde{V}$, $\xi \in \Upsilon_E^{\text{int}}$ is chosen as the MsFEM type hat functions as suggested in Section 3.2 above. However, it is independent on the mesh sizes h, H . A crucial ingredient to obtain this bound is the fact that we included this macroscopic partition of unity in our reduced approximation space \tilde{V} . If alternatively we would chose

$p_\xi \in \tilde{V}$ to be the traditional Lagrange hat functions, the bound on $c_{\text{pu}, \tilde{V}}$ in Proposition 5.7 would be independent of the contrast. In fact, we might expect that $c_{\text{pu}, \tilde{V}}$ behaves much better than the upper bound due to the approximation properties of the reduced space. It would actually be possible to compute $c_{\text{pu}, \tilde{V}}$ for given V_h, \tilde{V} which would however be computationally expensive and thus not of any use in practical applications. Proposition 5.7, however shows that the localized a posteriori error estimator in Corollary 5.3 in the context of ArbiLoMod is indeed robust and efficient, even with respect to $H \rightarrow 0$.

Comparing with other localized RB and multiscale methods, one observes a difference in the scaling of the efficiency constants. While in our case, c_{pu} is independent of both h and H , the a posteriori error estimator published for LRBMS has a H/h dependency [39, Theorem 4.6] and in the certification framework for SCRBE, a $h^{-1/2}$ scaling appears [43, Proposition 4.5]. The error estimators published for GMsFEM in [12] also have no dependency on H or h . However, they also rely on specific properties of the basis generation. Also in the analysis of the “Discontinuous Galerkin Reduced Basis Element Method” (DGRBE), Pacciarini et.al. have a factor of $h^{-1/2}$ in the a priori analysis [4] and in the a posteriori error estimator [40].

5.3 Local Efficiency

So far we did not use properties of the bilinear form other than coercivity and continuity. Assuming locality of the bilinear form as in (2.3), we get a local efficiency estimate and an improved global efficiency estimate.

Proposition 5.8 *Let the bilinear form a be given by (2.3). Then we have the localized efficiency estimate*

$$\|R_\mu(\tilde{u}_\mu)\|_{O'_\xi} \leq \gamma_\mu |u_\mu - \tilde{u}_\mu|_{\Omega_\xi, 1}. \quad (5.6)$$

Proof. Using the error identity

$$a_\mu(u_\mu - \tilde{u}_\mu, \varphi) = \langle R_\mu(\tilde{u}_\mu), \varphi \rangle,$$

we obtain for any $\varphi \in O_\xi$

$$\begin{aligned} \langle R_\mu(\tilde{u}_\mu), \varphi \rangle &= \int_{\Omega} \sigma_\mu(x) \nabla(u_\mu - \tilde{u}_\mu)(x) \nabla \varphi(x) dx \\ &= \int_{\Omega_\xi} \sigma_\mu(x) \nabla(u_\mu - \tilde{u}_\mu)(x) \nabla \varphi(x) dx \\ &\leq \gamma_\mu |u_\mu - \tilde{u}_\mu|_{\Omega_\xi, 1} \|\varphi\|_V, \end{aligned}$$

from which the statement follows. \square

Remark 5.9 *Under the assumptions of Proposition 5.8, it is easy to see that we have the improved efficiency estimate*

$$\Delta_{\text{loc}}(\tilde{u}_\mu) \leq \frac{\gamma_\mu \sqrt{c_{\text{ovlp}} c_{\text{pu}, \tilde{V}}}}{\alpha_\mu} \|u_\mu - \tilde{u}_\mu\|_V.$$

In many cases, a better constant can be found. Finite Element ansatz functions are usually not orthogonal if they share support. So if c_{ovlp} spaces have support in one point in space, they have to be placed in different groups when designing a partition for Proposition 5.2, so $c_{\text{ovlp}} \leq J$ (cf. [48, p. 67]).

5.4 Relative Error Bounds

From the error estimators for the absolute error, we can construct error estimators for the relative error. Estimates for the relative error are given in [23, Proposition 4.4], but the estimates used here are slightly sharper.

Proposition 5.10 *Assuming $\|\tilde{u}_\mu\|_V > \Delta(\tilde{u}_\mu)$ and $\|\tilde{u}_\mu\|_V > \Delta_{loc}(\tilde{u}_\mu)$, the error estimators defined by*

$$\begin{aligned}\Delta^{rel}(\tilde{u}_\mu) &:= \frac{\Delta(\tilde{u}_\mu)}{\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)} \\ \Delta_{loc}^{rel}(\tilde{u}_\mu) &:= \frac{\Delta_{loc}(\tilde{u}_\mu)}{\|\tilde{u}_\mu\|_V - \Delta_{loc}(\tilde{u}_\mu)}\end{aligned}$$

are robust and efficient:

$$\begin{aligned}\frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V} &\leq \Delta^{rel}(\tilde{u}_\mu) \leq \left(1 + 2\Delta^{rel}(\tilde{u}_\mu)\right) \frac{\gamma_\mu}{\alpha_\mu} \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V} \\ \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V} &\leq \Delta_{loc}^{rel}(\tilde{u}_\mu) \leq \left(1 + 2\Delta_{loc}^{rel}(\tilde{u}_\mu)\right) \frac{\gamma_\mu \sqrt{J} c_{pu, \tilde{V}}}{\alpha_\mu} \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V}\end{aligned}$$

Proof. Realizing that $(\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)) \leq \|u_\mu\|_V$, it is easy to see that

$$\frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V} \leq \frac{\Delta(\tilde{u}_\mu)}{\|u_\mu\|_V} \leq \frac{\Delta(\tilde{u}_\mu)}{\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)}, \quad (5.7)$$

which is the first inequality. Using

$$\|\tilde{u}_\mu\|_V + \Delta(\tilde{u}_\mu) = (\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)) \left(1 + 2\Delta^{rel}(\tilde{u}_\mu)\right)$$

the second inequality can be shown:

$$\begin{aligned}\Delta^{rel}(\tilde{u}_\mu) &= \frac{\Delta(\tilde{u}_\mu)}{\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)} \leq \frac{\gamma_\mu}{\alpha_\mu} \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|\tilde{u}_\mu\|_V - \Delta(\tilde{u}_\mu)} \\ &= \frac{\gamma_\mu}{\alpha_\mu} \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|\tilde{u}_\mu\|_V + \Delta(\tilde{u}_\mu)} \left(1 + 2\Delta^{rel}(\tilde{u}_\mu)\right) \\ &\leq \frac{\gamma_\mu}{\alpha_\mu} \frac{\|u_\mu - \tilde{u}_\mu\|_V}{\|u_\mu\|_V} \left(1 + 2\Delta^{rel}(\tilde{u}_\mu)\right).\end{aligned}$$

The inequalities for Δ_{loc}^{rel} can be shown accordingly. \square

Reviewing the five desired properties of an a posteriori error estimator at the beginning of this section, we see that the presented error estimator is robust and efficient (1) and is online-offline decomposable (2). Parallelization can be done over the spaces O_ξ . Only online data has to be transferred, so there is little communication (3). The online-offline decomposition only has to be repeated for a space O_ξ , if a new basis function with support in Ω_ξ was added. So reuse in unchanged regions is possible (4). How the adaptive enrichment is steered (5) will be described in the following section.

6 Enrichment Procedure

The first ArbiLoMod solution is obtained using the initial reduced local subspaces generated using the local training and greedy procedures described in Section 4. If this solution is not good enough according to the a posteriori error estimator, the solution is improved by enriching the reduced local subspaces and then solving the global reduced problem again. The full procedure is given in Algorithm 5 and described in the following.

For the enrichment, we use the overlapping local subspaces introduced in (5.5), which are also used for the a posteriori error estimator. Local problems are solved in the overlapping spaces. The original bilinear form is used, but as a right hand side the residual of the last reduced solution is employed. The local spaces and the parameter values for which the enrichment is performed are selected in a Dörfler-like [16] algorithm. The thus obtained local solutions u_l do not fit into our space decomposition, as they lie in one of the overlapping spaces, not in one of the local subspaces used for the basis construction. Therefore, the u_l are decomposed using the projection operators P_{V_ξ} defined in Definition 3.5. In the setting of our numerical example (Section 9), this decomposition yields at most 9 parts (one codim-2 part, four codim-1 parts and four codim-0 parts). Of these parts, the one worst approximated by the existing reduced local subspace is selected for enrichment. “Worst approximated” is here defined as having the largest part orthogonal to the existing reduced local subspace. We denote the part of $P_{V_\xi}(u_l)$ orthogonal to \tilde{V}_ξ w.r.t. the inner product of V by $(P_{V_\xi}(u_l))^\perp$.

To avoid communication, cell spaces $\tilde{V}_\xi, \xi \in \Upsilon_0$ are not enriched at this point. Such an enrichment would require the communication of the added basis vector, which might be large. Instead, only the other spaces are enriched, and the cell spaces associated with Υ_0 are regenerated using the greedy procedure from Section 4.3. For the other spaces, a strong compression of the basis vectors is possible (cf. Section 8).

This selection of the local spaces can lead to one reduced local space being enriched several times in one iteration. Numerical experiments have shown that this leads to poorly conditioned systems, as the enrichment might try to introduce the same feature into a local basis twice. To prevent this, the enrichment algorithm enriches each reduced local subspace at most once per iteration.

Algorithm 5: Online Enrichment

```

1 Function OnlineEnrichment( $d, \text{tol}$ ):
   Input : enrichment fraction  $d$ ,
           target error tol
2 while  $\max_{\mu \in \Xi} \Delta(\tilde{u}_\mu) > \text{tol}$  do
3      $E \leftarrow \emptyset$ 
4     while  $\left( \sum_{(\mu, \xi) \in E} \|R_\mu(\tilde{u}_\mu)\|_{(O_\xi)'} \right) / \left( \sum_{(\mu, \xi) \in (\Xi \times \Upsilon_E)} \|R_\mu(\tilde{u}_\mu)\|_{(O_\xi)'} \right) < d$  do
5          $\hat{\mu}, \hat{\xi} \leftarrow \arg \max_{(\mu, \xi) \in (\Xi \times \Upsilon_E) \setminus E} \|R_\mu(\tilde{u}_\mu)\|_{(O_\xi)'}$ 
6          $E \leftarrow E \cup (\hat{\mu}, \hat{\xi})$ 
7     /*  $S$  is used for double enrichment protection. */
8      $S \leftarrow \emptyset$ 
9     for  $(\mu, \xi) \in E$  do
10         find  $u_l \in O_\xi$  such that:
11              $a_\mu(u_l, \varphi) = \langle R_\mu(\tilde{u}_\mu), \varphi \rangle \quad \forall \varphi \in O_\xi$ 
12              $\check{\xi} \leftarrow \arg \max_{\xi \in \Upsilon \setminus \Upsilon_0} \|(P_{V_\xi}(u_l))^\perp\|_{V_\xi}$ 
13             if  $\check{\xi} \notin S$  then
14                  $\tilde{V}_{\check{\xi}} \leftarrow \tilde{V}_{\check{\xi}} \oplus \text{span}((P_{V_{\check{\xi}}}(u_l))^\perp)$ 
15                  $S \leftarrow S \cup \check{\xi}$ 
16     run LocalGreedy
17     recalculate reduced solutions
    
```

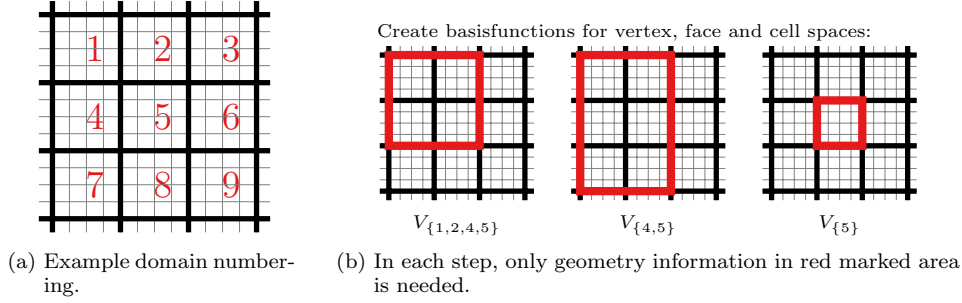


Figure 8.1: Before online enrichment, it is possible to compute all reduced basis function having support on a domain (Ω_5 here) using only local information about the domain and its surrounding domains.

7 Handling Local Changes

When the enrichment iteration has converged to a sufficient accuracy, the obtained solution is handed over to the user. The envisioned implementation should then wait for the user to modify the model under consideration. After a change to the model simulated, the full procedure described above is repeated, but wherever possible, existing data is reused. On the changed domains, a new mesh is generated, if necessary. Basis vectors having support in the changed region are discarded. The domain decomposition is never changed and thus has to be independent of the geometry. The potential savings are not only in the reduced basis generation, but also in the assembly of the system matrices. In an implementation featuring localized meshing and assembly domain by domain, meshing and assembly has to be repeated only in the domains affected by the change. This approach of recomputing everything in the changed region was chosen because it allows a robust implementation without any assumptions about the changes.

8 Runtime and Communication

A major design goal of ArbiLoMod is communication avoidance and scalability in parallel environments. Although the main topic of this publication are the mathematical properties, we want to highlight the possibilities offered by ArbiLoMod to reduce communication in a parallel setup.

Similar to overlapping Domain Decomposition methods we require, that not only the local domain, but also an overlap region is available locally. For a subdomain Ω_i the overlap region is the domain itself and all adjacent domains, as depicted in Figure 8.1, i.e. all subdomains in the neighborhood $\mathcal{N}_{\{i\}}$. As the overlap region includes the support of all training spaces, one can compute all initial reduced local subspaces with support in Ω_i without further communication. This work can be distributed on many nodes. Afterwards, only reduced representations of the operator have to be communicated. Using the operator decomposition $a^b(u, v) = \sum_{i=1}^{N_D} a_{\Omega_i}^b(u, v)$, a global, reduced operator is collected using an all-to-one communication of reduced matrices. The global reduced problem is then solved on a single node. It is assumed that the global, reduced system is sufficiently small.

If the accuracy is not sufficient, online enrichment is performed. This step requires additional communication; first for the evaluation of the error estimator and second to communicate new basis vectors of reduced face spaces \tilde{V}_ξ , $\xi \in \Upsilon_1$. Note that it is sufficient to communicate the local projection $P_{U_\xi}(\psi)$ and reconstruct the actual basis function as its extension, so that we save communication costs proportional to the volume to surface ratio.

The evaluation of the localized error estimator requires the dual norms of the residual




Figure 9.1: First structure in sequence of simulated structures. Unit square with high and low conductivity regions. White: constant conductivity region ($\sigma = 1$), black: parameterized high conductivity region ($\sigma_\mu = 1 + \mu$). Homogeneous Neumann boundaries Γ_N at top and bottom marked blue, inhomogeneous Dirichlet boundaries Γ_D at left and right marked light red.

in the localized spaces $\|R_\mu(\tilde{u}_\mu)\|_{O'_\xi}$. Using a stabilized online-offline splitting [8], the evaluation of the error estimator can be evaluated for the full system using only reduced quantities. The computation of the reduced operators is performed in parallel, similar to the basis construction in the first step. The actual evaluation of the error estimator can be performed on a single node and is independent of the number of degrees of freedom of the high fidelity model.

An important parameter for ArbiLoMod's runtime is the domain size H . The domain size affects the size of the local problems, the amount of parallelism in the algorithm, and the size of the reduced global problem. An H too large leads to large local problems, while an H too small leads to a large reduced global problem (see also the numerical example in Section 9.2 and especially the results in Table 9.3). H has to be chosen to balance these two effects. As the focus of this manuscript is ArbiLoMod's mathematical properties, the question of choosing H for optimal performance will be postponed to future research.

9 Numerical Example

The numerical experiments were performed using pyMOR [37]. The source code for the reproduction of all results presented in this section are provided as a supplement to this paper. See the README file therein for installation instructions. Note that this code is kept simple to easily explore ArbiLoMod's mathematical properties. It is not tuned for performance. First results for electrodynamics were published in [9].

9.1 Problem Definition

To illustrate the capabilities of ArbiLoMod we apply it to a sequence of locally modified geometries. We consider heat conduction without heat sources in the domain on the unit square $\Omega :=]0, 1[^2$. We approximate u solving

$$-\nabla \cdot (\sigma_\mu \nabla u_\mu) = 0 \quad (9.1)$$

where $\sigma_\mu : \Omega \rightarrow \mathbb{R}$ is the heat conductivity. We apply homogeneous Neumann boundaries at the top and the bottom: $\nabla u \cdot n = 0$ on $\Gamma_N := ([0, 1[\times 0) \cup ([0, 1[\times 1)$ and inhomogeneous Dirichlet boundaries at the left and right: $u = 1$ on $\Gamma_{D,1} := 0 \times]0, 1[$, $u = -1$ on $\Gamma_{D,-1} := 1 \times]0, 1[$ (see also Figure 9.1).

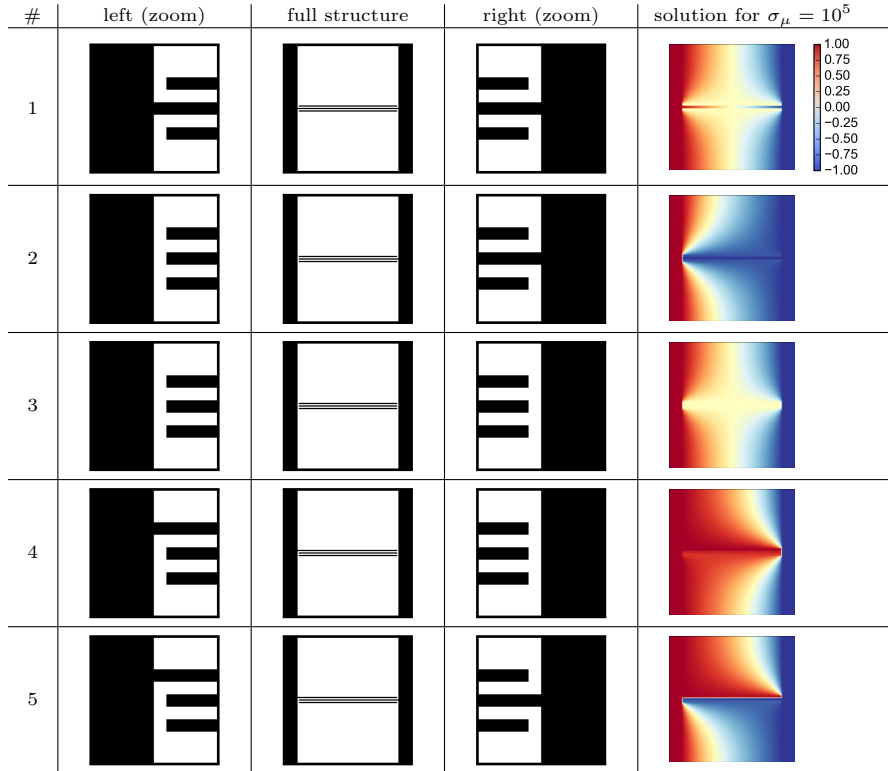


Figure 9.2: Sequence of structures simulated along with solutions for one parameter value. Very localized changes cause strong global changes in the solution. Script: `create_full_solutions.py`.

The unit square is partitioned into two regions: one region with constant heat conductivity $\sigma_\mu = 1$ and one region with constant, but parameterized conductivity $\sigma_\mu = 1 + \mu$ where $\mu \in [10^0, 10^5]$. We call this second region the “high conductivity region” $\Omega_{h,i}$. For reproduction or benchmarking, we give precise definitions in the following. In the first geometry, the high conductivity region is

$$\begin{aligned} \Omega_{h,1} := & \left[(0.0, 0.1) \times (0, 1) \right] \cup \left[(0.9, 1.0) \times (0, 1) \right] \cup \\ & \left[(0.11, 0.89) \times (0.475, 0.485) \right] \cup \left[(0.1, 0.9) \times (0.495, 0.505) \right] \cup \\ & \left[(0.11, 0.89) \times (0.515, 0.525) \right], \end{aligned} \quad (9.2)$$

see also Figure 9.1. The inhomogeneous Dirichlet boundary conditions are handled by a shift function u_s and we solve for u_0 having homogeneous Dirichlet values where $u = u_0 + u_s$. The parametrization of the conductivity in the high conductivity region, σ_μ , leads to a term in the affine decomposition of the bilinear form. The affine decomposition of the bilinear form and linear form are

$$\begin{aligned} a_\mu(u_0, \varphi) &= \mu \int_{\Omega_{h,i}} \nabla u_0 \cdot \nabla \varphi \, dx + \int_{\Omega} \nabla u_0 \cdot \nabla \varphi \, dx \\ \langle f_\mu, \varphi \rangle &= -\mu \int_{\Omega_{h,i}} \nabla u_s \cdot \nabla \varphi \, dx - \int_{\Omega} \nabla u_s \cdot \nabla \varphi \, dx. \end{aligned} \quad (9.3)$$

The coercivity constant α_μ of the corresponding bilinear form with respect to the H^1 norm is bounded from below by $\alpha_{LB} := \frac{\sigma_{\min}}{c_F^2 + 1}$, where σ_{\min} is the minimal conductivity and $c_F = \frac{1}{\sqrt{2\pi}}$ is the constant in the Friedrich’s inequality $\|\varphi\|_{L^2(\Omega)} \leq c_F \|\nabla \varphi\|_{L^2(\Omega)} \forall \varphi \in H_0^1(\Omega)$. The problem is discretized using P^1 ansatz functions on a structured triangle grid with maximum triangle size h . The grid is carefully constructed to resolve the high conductivity regions, i.e. h is chosen to be $1/n$ where n is a multiple of 200. To mimic “arbitrary local modifications”, the high conductivity region is changed slightly four times, which leads to a sequence of five structures to be simulated in total. The high conductivity regions are defined as:

$$\begin{aligned} \Omega_{h,2} &:= \Omega_{h,1} \setminus \left[(0.1, 0.11) \times (0.495, 0.505) \right] \\ \Omega_{h,3} &:= \Omega_{h,2} \setminus \left[(0.89, 0.9) \times (0.495, 0.505) \right] \\ \Omega_{h,4} &:= \Omega_{h,3} \cup \left[(0.1, 0.11) \times (0.515, 0.525) \right] \\ \Omega_{h,5} &:= \Omega_{h,4} \cup \left[(0.89, 0.9) \times (0.495, 0.505) \right] \end{aligned} \quad (9.4)$$

These modifications only affect a very small portion of the domain (actually, only 0.01%), but for high contrast configurations, they lead to strong global changes in the solution, see Figure 9.2. An equidistant domain decomposition of 8×8 domains is used. The mesh resolves the domain boundaries.

Configuration

If not specified otherwise, we use a mesh size of $1/h = 200$, a training tolerance of $\varepsilon_{\text{train}} = 10^{-4}$, a number of random samplings of $M = 60$, a greedy tolerance of $\varepsilon_{\text{greedy}} = 10^{-3}$, a convergence criterion of $\|R_\mu(\tilde{u}_\mu)\|_{V_h'} < 10^{-2}$, an enrichment fraction of $d = 0.5$, a training set of size $|\Xi| = 6$, and the parameter for extension calculation is $\bar{\mu} = 10^5$.

9.2 Results

The initial reduced space is created using the local trainings and greedy algorithms. In both the trainings and the greedy algorithms a tolerance parameter steers the quality of the obtained reduced space: In the trainings, $\varepsilon_{\text{train}}$ is the stopping criterion for the




Figure 9.3: Maximum relative H^1 -error on training set Ξ in dependence of tolerances in codim 1 training and codim 0 greedy. Online enrichment disabled. Script: `experiment_tolerances.py`.

geometry	trainings		with training greedys		iterations		without training iterations	
	reuse:		reuse:		reuse:		reuse:	
	no	yes	no	yes	no	yes	no	yes
1	112	112 (-0 %)	64	64 (-0 %)	24	24 (-0 %)	46	46 (-0 %)
2	112	5 (-96 %)	64	8 (-88 %)	24	13 (-46 %)	48	28 (-42 %)
3	112	5 (-96 %)	64	8 (-88 %)	20	14 (-30 %)	42	27 (-36 %)
4	112	3 (-97 %)	64	6 (-91 %)	25	10 (-60 %)	54	23 (-57 %)
5	112	5 (-96 %)	64	8 (-88 %)	25	12 (-52 %)	52	27 (-48 %)

Table 9.1: Number of iterations of online enrichment: (a) With and without codim 1 training. (b) With and without reuse of basis functions of previous simulations. Convergence criterion: $\|R_\mu(\tilde{u}_\mu)\|_{V'_h} < 10^{-4}$, greedy tolerance: $\varepsilon_{\text{greedy}} = 10^{-5}$. See also Figures 9.4, 9.5.

Scripts: `experiment_basisreuse_with_training.py`,
`experiment_basisreuse.py`.

SnapshotGreedy (Algorithm 2) and the local greedys stop when the local error estimator stays below the prescribed tolerance $\varepsilon_{\text{greedy}}$. The resulting reduction errors in dependence on the two tolerances are depicted in Figure 9.3.

If the resulting error is too big, it can be further reduced using iterations of online enrichment as depicted in Figure 9.4. Results suggest an very rapid decay of the error with online enrichment. The benefits of ArbiLoMod can be seen in Figure 9.4 and Table 9.1: after the localized geometry changes, most of the work required in the initial basis creation does not need to be repeated and the online enrichments converge faster for subsequent simulations, leading to less iterations. The online enrichment presented here converges even when started on empty bases, as depicted in Figure 9.5. It does not rely on properties of the reduced local subspaces created by trainings and greedys. The performance of the localized a posteriori error estimator Δ_{loc}^{rel} can be seen in Figure 9.6. Comparison of the localized estimator Δ_{loc}^{rel} with the global estimator Δ^{rel} shows that, for the example considered here, the localization does not add a significant factor beyond the factor $c_{\text{pu}, \tilde{V}}$, which is supposed to be close to one and was thus neglected in Fig. 9.6.

Even though our implementation was not tuned for performance and is not parallel, we present some timing measurements in Table 9.2 and Table 9.3. Table 9.2 shows that, already in our unoptimized implementation, trainings and greedys have a shorter runtime than a single global solve for problems of sufficient size. Taking into account that trainings and greedys create a solution space valid in the whole parameter space, this data is a strong hint that ArbiLoMod can realize its potential for acceleration for large problems, in an




Figure 9.4: Relative error over iteration with and without basis reuse after geometry change. With codim 1 training. Convergence criterion: $\|R_\mu(\tilde{u}_\mu)\|_{V'_h} < 10^{-4}$, greedy tolerance: $\varepsilon_{\text{greedy}} = 10^{-5}$. See also Table 9.1. Script: `experiment.basisreuse.with.training.py`.




Figure 9.5: Relative H^1 error over iterations with and without basis reuse after geometry change. Without codim 1 training. Convergence criterion: $\|R_\mu(\tilde{u}_\mu)\|_{V'_h} < 10^{-4}$, greedy tolerance: $\varepsilon_{\text{greedy}} = 10^{-5}$. See also Table 9.1. Script: `experiment.basisreuse.py`.




Figure 9.6: Error estimators Δ^{rel} and Δ_{loc}^{rel} over iterations, compared to the relative error. Plotted for $\alpha_\mu = c_{\text{pu}, \tilde{V}} = 1$. Simulation performed with a convergence criterion of $\|R_\mu(\tilde{u}_\mu)\|_{V'_h} < 10^{-6}$, a training tolerance of $\varepsilon_{\text{train}} = 10^{-5}$, and a greedy tolerance of $\varepsilon_{\text{greedy}} = 10^{-7}$. Script: `experiment_estimatorperformance.py`.

$1/h$	global dofs	global solve time [s]	# dofs, codim 1 training space	avg. time per codim 1 training [s]	# dofs, codim 0 space	max time per codim 0 greedy [s]	# dofs, reduced problem	solve time, reduced [ms]	max error [% ₀]
200	80,401	0.656	7,626	1.02	1,201	4.9	1,178	21.8	1.316
400	320,801	4.87	30,251	5.14	4,901	7.04	1,151	22.4	1.433
600	721,201	23.6	67,876	14	11,101	10.7	1,116	19.1	2.035
800	1,281,601	41.8	120,501	29.5	19,801	17.8	1,101	17.1	2.735
1000	2,002,001	86.4	188,126	51.3	31,001	24.2	1,089	18.8	1.351
1200	2,882,401	230	270,751	81.2	44,701	36.6	1,082	18.6	4.462
1400	3,922,801	230	368,376	120	60,901	51.7	1,073	18.2	2.379

Table 9.2: Runtimes for selected parts of ArbiLoMod without online enrichment. “max error” denotes $\max_{\mu \in \Xi} \|u_\mu - \tilde{u}_\mu\|_V / \|u_\mu\|_V$. Runtimes measured using a pure Python implementation, using SciPy solvers (SuperLU sequential). Note that the global solve time is for one parameter value while training and greedy produce spaces valid for all parameter values in the training set Ξ . Script: `experiment_create_timings.py`.

$1/H$	# dofs, codim 1 training space	mean training time [s]	max training time [s]	# dofs, codim 0 space	mean greedy time [s]	max greedy time [s]	# dofs, reduced problem	solve time, reduced problem	max error [%]
4	30,251	10.9	14.1	4,901	3.07	6.6	403	3.85	0.362
5	19,401	6.88	8.74	3,121	2.32	11.7	517	4.38	0.435
8	7,626	2.63	3.55	1,201	1.35	5.69	1,178	14.9	1.32
10	4,901	1.68	1.94	761	0.655	3.98	1,451	19.8	0.220
20	1,251	0.483	0.661	181	0.305	3.18	5,025	94.4	0.0804

Table 9.3: Influence of domain size H . Fine mesh resolution: $1/h = 200$. “max error” is $\max_{\mu \in \Xi} \|u_\mu - \tilde{u}_\mu\|_V / \|u_\mu\|_V$. Smaller domains lead to more parallelism and smaller local problems, but also to more global dofs and a worse constant in the a-posteriori error estimator. Measured using a pure Python implementation, using SciPy solvers (SuperLU sequential). Script: `experiment_H_study.py`.




Figure 9.7: Distribution of local basis sizes after initial training. Relative reduction error at this configuration: $1.3 \cdot 10^{-3}$. Script: `experiment_draw_basis_sizes.py`.

optimized implementation and a parallel computing environment. Especially when the solution is to be calculated at multiple parameter values. Table 9.3 shows the effect of choosing the domain size H : Large domains lead to large local problems, while small domains lead to a large reduced global problem (see also Section 8).

10 Conclusion

We introduced ArbiLoMod, a simulation technique aiming at problems with arbitrary local modifications and highly parallel computing environments. It is based on the Reduced Basis method, inheriting its advantages, but localizing the basis construction and error estimation. It consists of basis generation algorithms, a localized a posteriori error estimator controlling the reduction error, and a localized space enrichment procedure, improving the reduced local spaces if necessary. The initial basis generation algorithms require no communication of unreduced quantities in a parallel implementation. For the basis enrichment procedure, a strong compression of the communicated quantities is possible. We discussed the possibilities to use ArbiLoMod to implement a scalable parallel code by reducing communication costs due to the local structure and the local Reduced Basis strategy. ArbiLoMod was demonstrated on a coercive example in two dimensions, featuring high contrast, fine details and channels. Even though small local modifications to this example lead to strong global changes in the solution, ArbiLoMod was able to approximate the new solutions after these geometry changes with a small fraction of the effort needed for the initial geometry. BSD-licensed source code is provided along with

this publication, anyone can reproduce all results presented here easily.

Acknowledgment

The authors would like to thank Clemens Pechstein for many fruitful discussions and the referees, whose reviews lead to significant improvements of this manuscript.

References

- [1] A. Abdulle and Y. Bai. Reduced basis finite element heterogeneous multiscale method for high-order discretizations of elliptic homogenization problems. *Journal of Computational Physics*, 231(21):7014 – 7036, 2012.
- [2] A. Abdulle and Y. Bai. Adaptive reduced basis finite element heterogeneous multiscale method. *Computer Methods in Applied Mechanics and Engineering*, 257:203 – 220, 2013.
- [3] F. Albrecht and M. Ohlberger. The localized reduced basis multi-scale method with online enrichment. *Oberwolfach Reports*, vol.7:406–409, 2013.
- [4] P. F. Antonietti, P. Pacciarini, and A. Quarteroni. A discontinuous Galerkin reduced basis element method for elliptic problems. *ESAIM: M2AN*, doi 10.1051/m2an/2015045, Jun 2015.
- [5] I. Babuška, X. Huang, and R. Lipton. Machine computation using the exponentially convergent multiscale spectral generalized finite element method. *ESAIM: M2AN*, 48(02):493–515, 2014.
- [6] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk. Convergence rates for greedy algorithms in reduced basis methods. *SIAM J. Math. Anal.*, 43(3):1457–1472, Jan 2011.
- [7] A. Buhr. Finite elements in PCB structures. Master’s thesis, TU Darmstadt, 2009.
- [8] A. Buhr, C. Engwer, M. Ohlberger, and S. Rave. A numerically stable a posteriori error estimator for reduced basis approximations of elliptic equations. In X. O. E. Onate and A. Huerta, editors, *Proceedings of the 11th World Congress on Computational Mechanics*, pages 4094–4102. CIMNE, Barcelona, 2014.
- [9] A. Buhr, C. Engwer, M. Ohlberger, and S. Rave. Arbilomod: Local solution spaces by random training in electrodynamics. *arXiv e-prints*, (1606.06206), june 2016. <https://arxiv.org/abs/1606.06206>.
- [10] A. Buhr and M. Ohlberger. Interactive simulations using localized reduced basis methods. In *IFAC-PapersOnLine*, volume 48(1), pages 729–730, 2015. 8th Vienna International Conference on Mathematical Modelling MATHMOD 2015.
- [11] Y. Chen, J. S. Hesthaven, Y. Maday, and J. Rodriguez. A monotonic evaluation of lower bounds for inf-sup stability constants in the frame of reduced basis approximations. *Comptes Rendus Mathematique*, 346(23):1295 – 1300, 2008.
- [12] E. T. Chung, Y. Efendiev, and W. T. Leung. An adaptive generalized multiscale discontinuous Galerkin method (GMSDGM) for high-contrast flow problems. *arXiv preprint arXiv:1409.3474*, Sept. 2014.
- [13] E. T. Chung, Y. Efendiev, and W. T. Leung. Residual-driven online generalized multiscale finite element methods. *arXiv preprint arXiv:1501.04565*, 2015.

- [14] E. T. Chung, Y. Efendiev, and G. Li. An adaptive GMsFEM for high-contrast flow problems. *Journal of Computational Physics*, 273:54–76, Sep 2014.
- [15] A. Cohen, R. DeVore, and R. H. Nochetto. Convergence rates of AFEM with H^{-1} data. *Found. Comput. Math.*, 12(5):671–718, 2012.
- [16] W. Dörfler. A convergent adaptive algorithm for poisson’s equation. *SIAM Journal on Numerical Analysis*, 33(3):1106–1124, 1996.
- [17] W. E and B. Engquist. The heterogeneous multi-scale methods. *Comm. Math. Sci*, 1:87–132, 2003.
- [18] Y. Efendiev, J. Galvis, and T. Y. Hou. Generalized multiscale finite element methods (GMsFEM). *Journal of Computational Physics*, Jan. 2013.
- [19] J. L. Eftang and A. T. Patera. Port reduction in parametrized component static condensation: approximation and a posteriori error estimation. *International Journal for Numerical Methods in Engineering*, 96(5):269–302, Jul 2013.
- [20] J. L. Eftang and A. T. Patera. A port-reduced static condensation reduced basis element method for large component-synthesized structures: approximation and a posteriori error estimation. *Advanced Modeling and Simulation in Engineering Sciences*, 1:3, 2014.
- [21] B. Haasdonk. Reduced basis methods for parametrized PDEs – A tutorial introduction for stationary and instationary problems. Technical report, 2014. Chapter to appear in P. Benner, A. Cohen, M. Ohlberger and K. Willcox: ”Model Reduction and Approximation: Theory and Algorithms”, SIAM.
- [22] P. Henning and D. Peterseim. Oversampling for the multiscale finite element method. *Multiscale Modeling & Simulation*, 11(4):1149–1175, 2013.
- [23] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. SpringerBriefs in Mathematics. Springer International Publishing, 2016.
- [24] J. S. Hesthaven, S. Zhang, and X. Zhu. Reduced basis multiscale finite element methods for elliptic problems. *Multiscale Modeling & Simulation*, 13(1):316–337, 2015.
- [25] U. Hetmaniuk and A. Klawonn. Error estimates for a two-dimensional special finite element method based on component mode synthesis. *Electronic Transactions on Numerical Analysis*, 41:109–132, 2014.
- [26] U. L. Hetmaniuk and R. B. Lehoucq. A special finite element method based on component mode synthesis. *ESAIM: M2AN*, 44(3):401–420, Feb 2010.
- [27] T. Y. Hou and X.-H. Wu. A multiscale finite element method for elliptic problems in composite materials and porous media. *Journal of Computational Physics*, 134:169–189, 1997.
- [28] T. J. Hughes, G. R. Feijo, L. Mazzei, and J.-B. Quincy. The variational multiscale methoda paradigm for computational mechanics. *Computer Methods in Applied Mechanics and Engineering*, 166(1):3 – 24, 1998.
- [29] D. Huynh, G. Rozza, S. Sen, and A. Patera. A successive constraint linear optimization method for lower bounds of parametric coercivity and infsup stability constants. *Comptes Rendus Mathematique*, 345(8):473 – 478, 2007.

- [30] L. Iapichino. *Reduced basis methods for the solution of parametrized PDEs in repetitive and complex networks with application to CFD*. PhD thesis, 2012.
- [31] L. Iapichino, A. Quarteroni, and G. Rozza. A reduced basis hybrid method for the coupling of parametrized domains represented by fluidic networks. *Computer Methods in Applied Mechanics and Engineering*, 221-222:63–82, May 2012.
- [32] A. Klawonn and O. B. Widlund. Dual-primal FETI methods for linear elasticity. *Communications on pure and applied mathematics*, 59(11):1523–1572, 2006.
- [33] Y. Maday and E. M. Rønquist. A reduced-basis element method. *Journal of Scientific Computing*, 17(1/4):447–459, 2002.
- [34] Y. Maday and E. M. Ronquist. The reduced basis element method: Application to a thermal fin problem. *SIAM Journal on Scientific Computing*, 26(1):240–258, Jan 2004.
- [35] I. Maier and B. Haasdonk. A Dirichlet-Neumann reduced basis method for homogeneous domain decomposition problems. *Applied Numerical Mathematics*, 78:31–48, 2014.
- [36] I. Martini, G. Rozza, and B. Haasdonk. Reduced basis approximation and a-posteriori error estimation for the coupled stokes-darcy system. *Advances in Computational Mathematics*, pages 1–27, 2014.
- [37] R. Milk, S. Rave, and F. Schindler. pyMOR-generic algorithms and interfaces for model order reduction. *arXiv preprint arXiv:1506.07094*, 2015.
- [38] A. Mlqvist and D. Peterseim. Localization of elliptic multiscale problems. *Math. Comp.*, 83(290):25832603, Jun 2014.
- [39] M. Ohlberger and F. Schindler. Error control for the localized reduced basis multiscale method with adaptive on-line enrichment. *SIAM Journal on Scientific Computing*, 37(6):A2865–A2895, 2015.
- [40] P. Pacciarini. *Discontinuous Galerkin Reduced Basis Element Methods for Parametrized Partial Differential Equations in Partitioned Domains*. PhD thesis, Politecnico Di Milano, 2016.
- [41] D. B. Phuong Huynh, D. J. Knezevic, and A. T. Patera. A static condensation reduced basis element method : approximation and a posteriori error estimation. *M2AN*, 47(1):213–251, Nov 2012.
- [42] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced Basis Methods for Partial Differential Equations*. La Matematica per il 3+2. Springer International Publishing, 2016.
- [43] K. Smetana. A new certification framework for the port reduced static condensation reduced basis element method. *Computer Methods in Applied Mechanics and Engineering*, 283:352–383, Jan 2015.
- [44] K. Smetana and A. T. Patera. Optimal local approximation spaces for component-based static condensation procedures. *Technical Report, MIT, No. 70*, 2015.
- [45] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl. Abstract robust coarse spaces for systems of pdes via generalized eigenproblems in the overlaps. *Numerische Mathematik*, 126(4):741–770, Aug 2013.

- [46] T. Strouboulis, I. Babuška, and K. Copps. The design and analysis of the generalized finite element method. *Computer methods in applied mechanics and engineering*, 181(1):43–69, 2000.
- [47] T. Strouboulis, K. Copps, and I. Babuška. The generalized finite element method. *Computer methods in applied mechanics and engineering*, 190(32):4081–4193, 2001.
- [48] A. Toselli and O. Widlund. *Domain decomposition methods—algorithms and theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 2005.